


CONNEXIONS



The Interoperability Report

December 1988

Volume 2, No. 12

*ConneXions —
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

Transition and coexistence for TCP/IP to OSI.....	2
The new NSFNET backbone..	6
Background File Transfer....	10
Directory of computer nets...	14
Looking ahead to 1989.....	15

ConneXions is published by Advanced Computing Environments, 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. Phone: 415-941-3399.

© 1988
Advanced Computing Environments.
Quotation with attribution encouraged.

ConneXions—The Interoperability Report
and the *ConneXions* masthead are
trademarks of Advanced Computing
Environments.

ISSN 0894-5926

From the Editor

This issue is being released at the *OSI Product Integration Conference*. Our feature article is an introduction to one of the tutorials in the conference. The article is entitled "Transition and Coexistence for TCP/IP to OSI" and is by Marshall Rose of The Wollongong Group. Given the large installed base of TCP/IP systems in the world, it is important to consider how to make a smooth transition to OSI protocols. Dr. Rose outlines two approaches in this article, and will discuss an additional three in the tutorial.

In our June 1987 issue [Volume 1, No. 2] we had an article on NSFNET which is one of the larger components of the Internet. Since then, a number of changes have taken place, most notably an upgrade of the backbone from 56kbps to T1 (1.544Mbps) links. Hans-Werner Braun of Merit Computer Network and The University of Michigan describes the new architecture in the first of a series of articles on the new NSFNET.

The second most popular application on TCP/IP networks is the File Transfer Protocol (FTP). (The most popular is of course electronic mail). Until now, FTP programs have been somewhat tedious to use due to their interactive nature which forces users to wait while their files are being transferred across often slow network links. A new system, called The Background File Transfer Program (BFTP), could solve many of FTP's problems. The designers of BFTP, Annette DeSchon and Bob Braden of USC-ISI, give a description of the system in an article on page 10.

The University of Texas System Office of Telecommunication Services publishes a "*Users' Directory of Computer Networks Accessible to the Texas Higher Education Network Member Institutions*." This directory is also applicable to other users of computer networks and serves as a useful reference guide. Tracy LaQuey from the University of Texas gives an overview of the directory on page 14.

Finally, on page 15, we look ahead to some upcoming articles and events for 1989. Have a nice holiday and see you in 1989!

Transition and Coexistence for TCP/IP to OSI

by Marshall T. Rose, The Wollongong Group, Inc.

One of the tutorials at the OSI Product Integration Conference is entitled "Issues in Transition and Coexistence for TCP/IP to OSI." This article serves as an introduction to the topics presented in the tutorial.

Introduction

The U.S. DoD Internet suite of protocols (commonly known as TCP/IP) is the de facto open (non-proprietary) standard for computer-communications in both multi-vendor and multi-administration networks. TCP/IP has enjoyed unprecedented success as the open systems solution of choice for inter-connecting networks and hosts.

However, based on international cooperative work, it is commonly acknowledged that protocols based on the Open Systems Interconnection (OSI) model and promulgated by the International Organization for Standardization (IEC/ISO) will eventually achieve dominance and enjoy even greater success than TCP/IP.

Although previously an "academic" problem, the widespread investment in TCP/IP-based systems has made practical solutions to transition and coexistence an overwhelming concern: organizations using TCP/IP protocols today will be less willing to adopt OSI protocols tomorrow unless interruption of production facilities is minimized and the underlying investment is protected. Considering the large installed base of TCP/IP today, and the continued growth of TCP/IP, there will be a tremendous problem when the time to move to OSI finally arrives!

The purpose of the tutorial is to introduce and compare different approaches to the transition problem. However, it is beyond the scope of this article to either define or contrast the Internet and OSI protocol suites. (Ed.: See [Comer] and [Tannenbaum]).

Metrics for comparison

If we are going to make comparisons, then we need to establish a checklist of good and bad features among the transition approaches. For our purposes, there are four sets of questions to ask:

- *Performance*: How well does the approach perform in terms of both throughput, latency, and host processing overhead? How does the approach impact the performance of other applications running in the network?
- *Flexibility*: What is the range of applicability of the approach? Is a special-purpose system required for each application, or can one general-purpose system serve the needs of a wide range of applications?
- *Transparency*: Is it possible for end-users to be unaware that the coexistence/transition approach is "in the loop?"
- *Pervasiveness*: How manageable is the approach? Does the approach impose additional administrative burden on the network operators? Do users' hosts (end-systems) need to be modified? Do internetwork relay hosts (intermediate systems) need to be modified?

All answers to these questions are necessarily subjective, although performance characteristics may be objectively compared given a sufficiently rigorous set of benchmarking definitions.

In the tutorial, we focus on five different transition technologies, but, owing to space limitations, this article will introduce only two.

Application Gateways

The application-gateway approach is a well-known, but often little understood, technology used to achieve interoperability between similar applications from different protocol suites. The most common use of this approach is for store-and-forward applications such as message handling. For instance, gateways between the Arpanet mail system and other mail systems have been in existence for quite some time. However, for end-to-end (non-store-and-forward) applications such as FTP and Telnet, this approach usually performs too poorly to be effective.

Consider the high-level architecture of an application-gateway host as shown in Figure 1. It is important to emphasize that this approach joins together two different application protocols and applies some sort of translation mechanism between the two. If we were interested in electronic mail, then APPL- α would be X.400, the OSI message handling service (MHS), while APPL- γ might be the Simple Mail Transfer Protocol (SMTP), the message transfer protocol in the Internet suite.

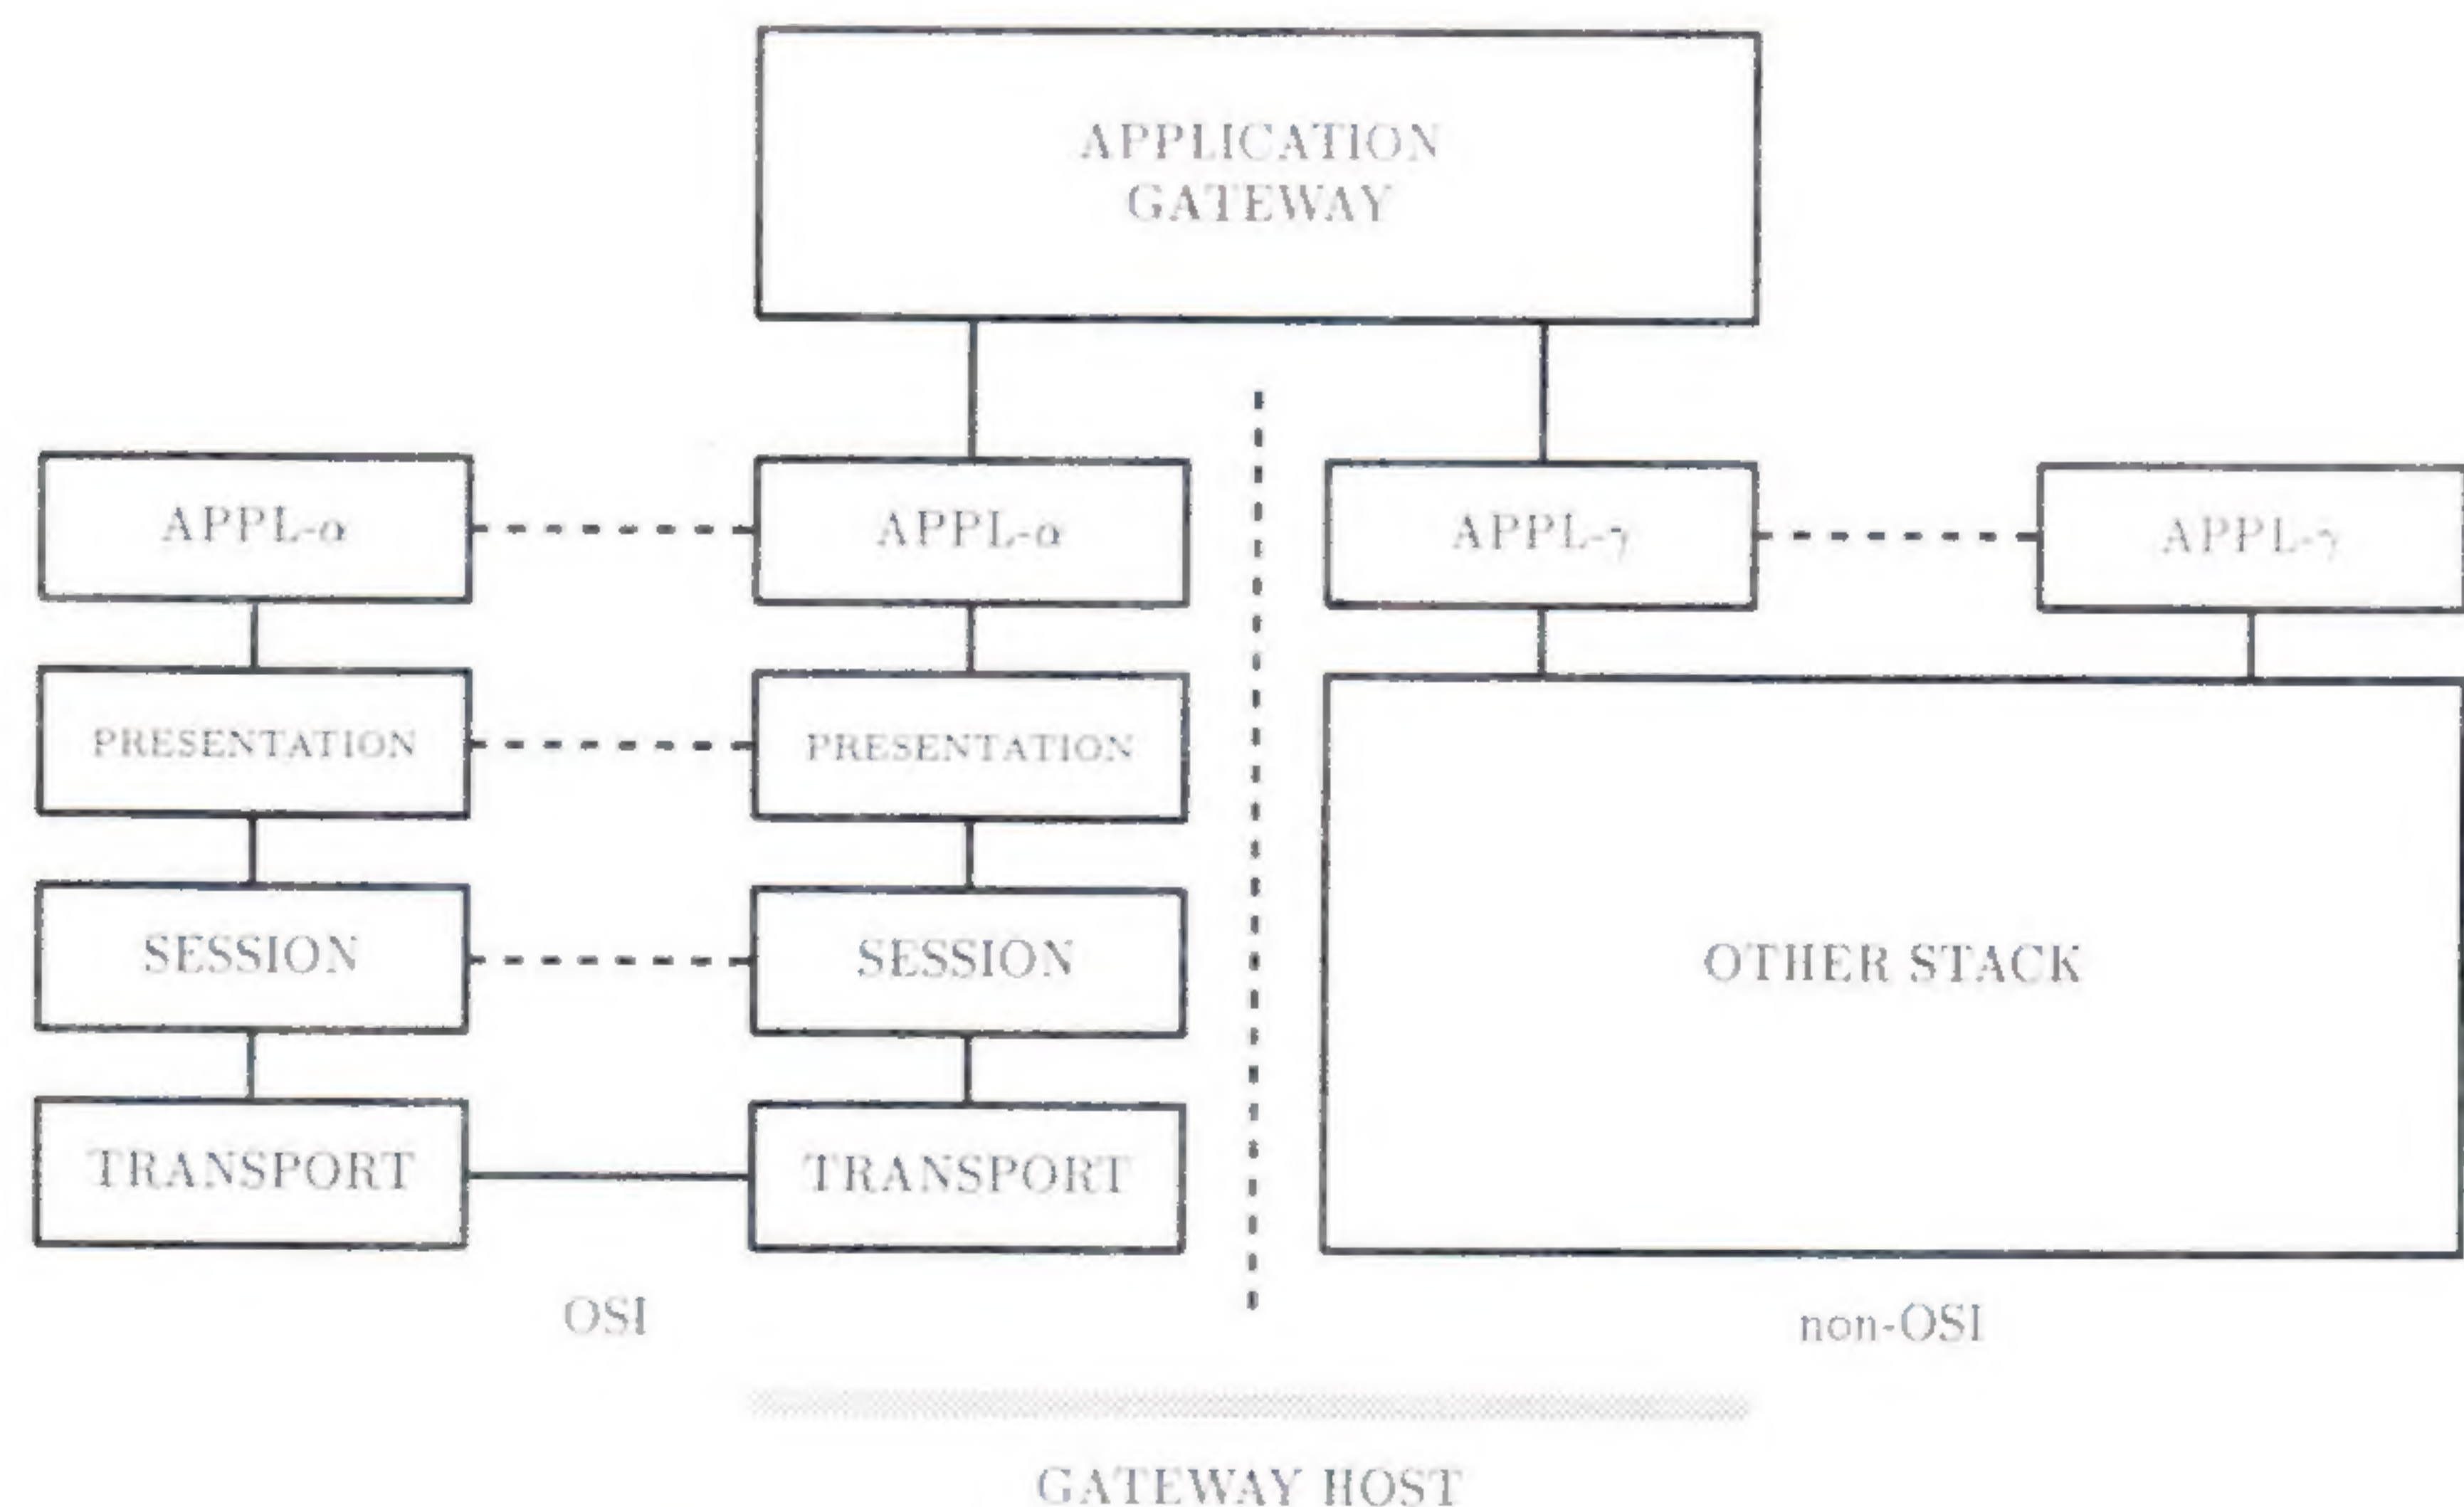


Figure 1: Application Gateway

Both are applications performing electronic mail store-and-forward functions, but the services offered by the protocols differ a great deal. The differences come from divergent views as to what "electronic mail" means. At the coarsest level of distinction, MHS offers a multi-media mail facility, whereas SMTP (together with the standard describing the format of messages, RFC 822) permits the exchange of text-oriented messages. The services offered by MHS are wide and varied, whereas SMTP offers a simple, basic store-and-forward capability, as its name suggests. Because of this, any translation between the two "electronic mail" protocols *must* result in a loss of information. As a result, a message that is sent through the gateway first in one direction and then the other will result in a different message than the original!

Transition/coexistence for TCP/IP to OSI (*continued*)

Because of these problems, application-gateways are said to solve *connectivity* problems but are poor solutions for *interoperability* problems.

Scorecard

So, how do application gateways “stack up” (no pun intended) against the metrics defined earlier?

- *Performance*: Acceptable for store-and-forward application classes.
- *Flexibility*: None, if a new application is to be supported, a new application-gateway must be built.
- *Transparency*: with respect to:
 - service*: Significant functionality is often lost due to inability to perform invertible mappings of services offered.
 - users*: Possible to achieve transparency; however, user is often required to apply “out-of-band” knowledge, e.g., embed a second host name (somehow) in the connection command.
- *Pervasiveness*: No end-system modifications are required. However, management of the application-gateways introduces administrative problems for end-to-end application classes, e.g., authentication.

Network-Service Tunnels

An alternate transition approach is to use a Network-Service Tunnel or *NS-Tunnel*. The idea is simple: CLNP (the OSIified version of IP) is encapsulated inside of the DoD IP, treating IP as simply a data-link protocol. As shown in Figure 2, the NS-Tunnel will perform as a router, conceptually removing one data-link header and adding another. This approach will require common protocols at the transport layer and above on both end-systems. However it does not require all intervening routers to use the same network protocol. This is an important feature: it implies that, with careful planning, transition of the network backbone may occur in several phases rather than changing everything in one fell swoop.

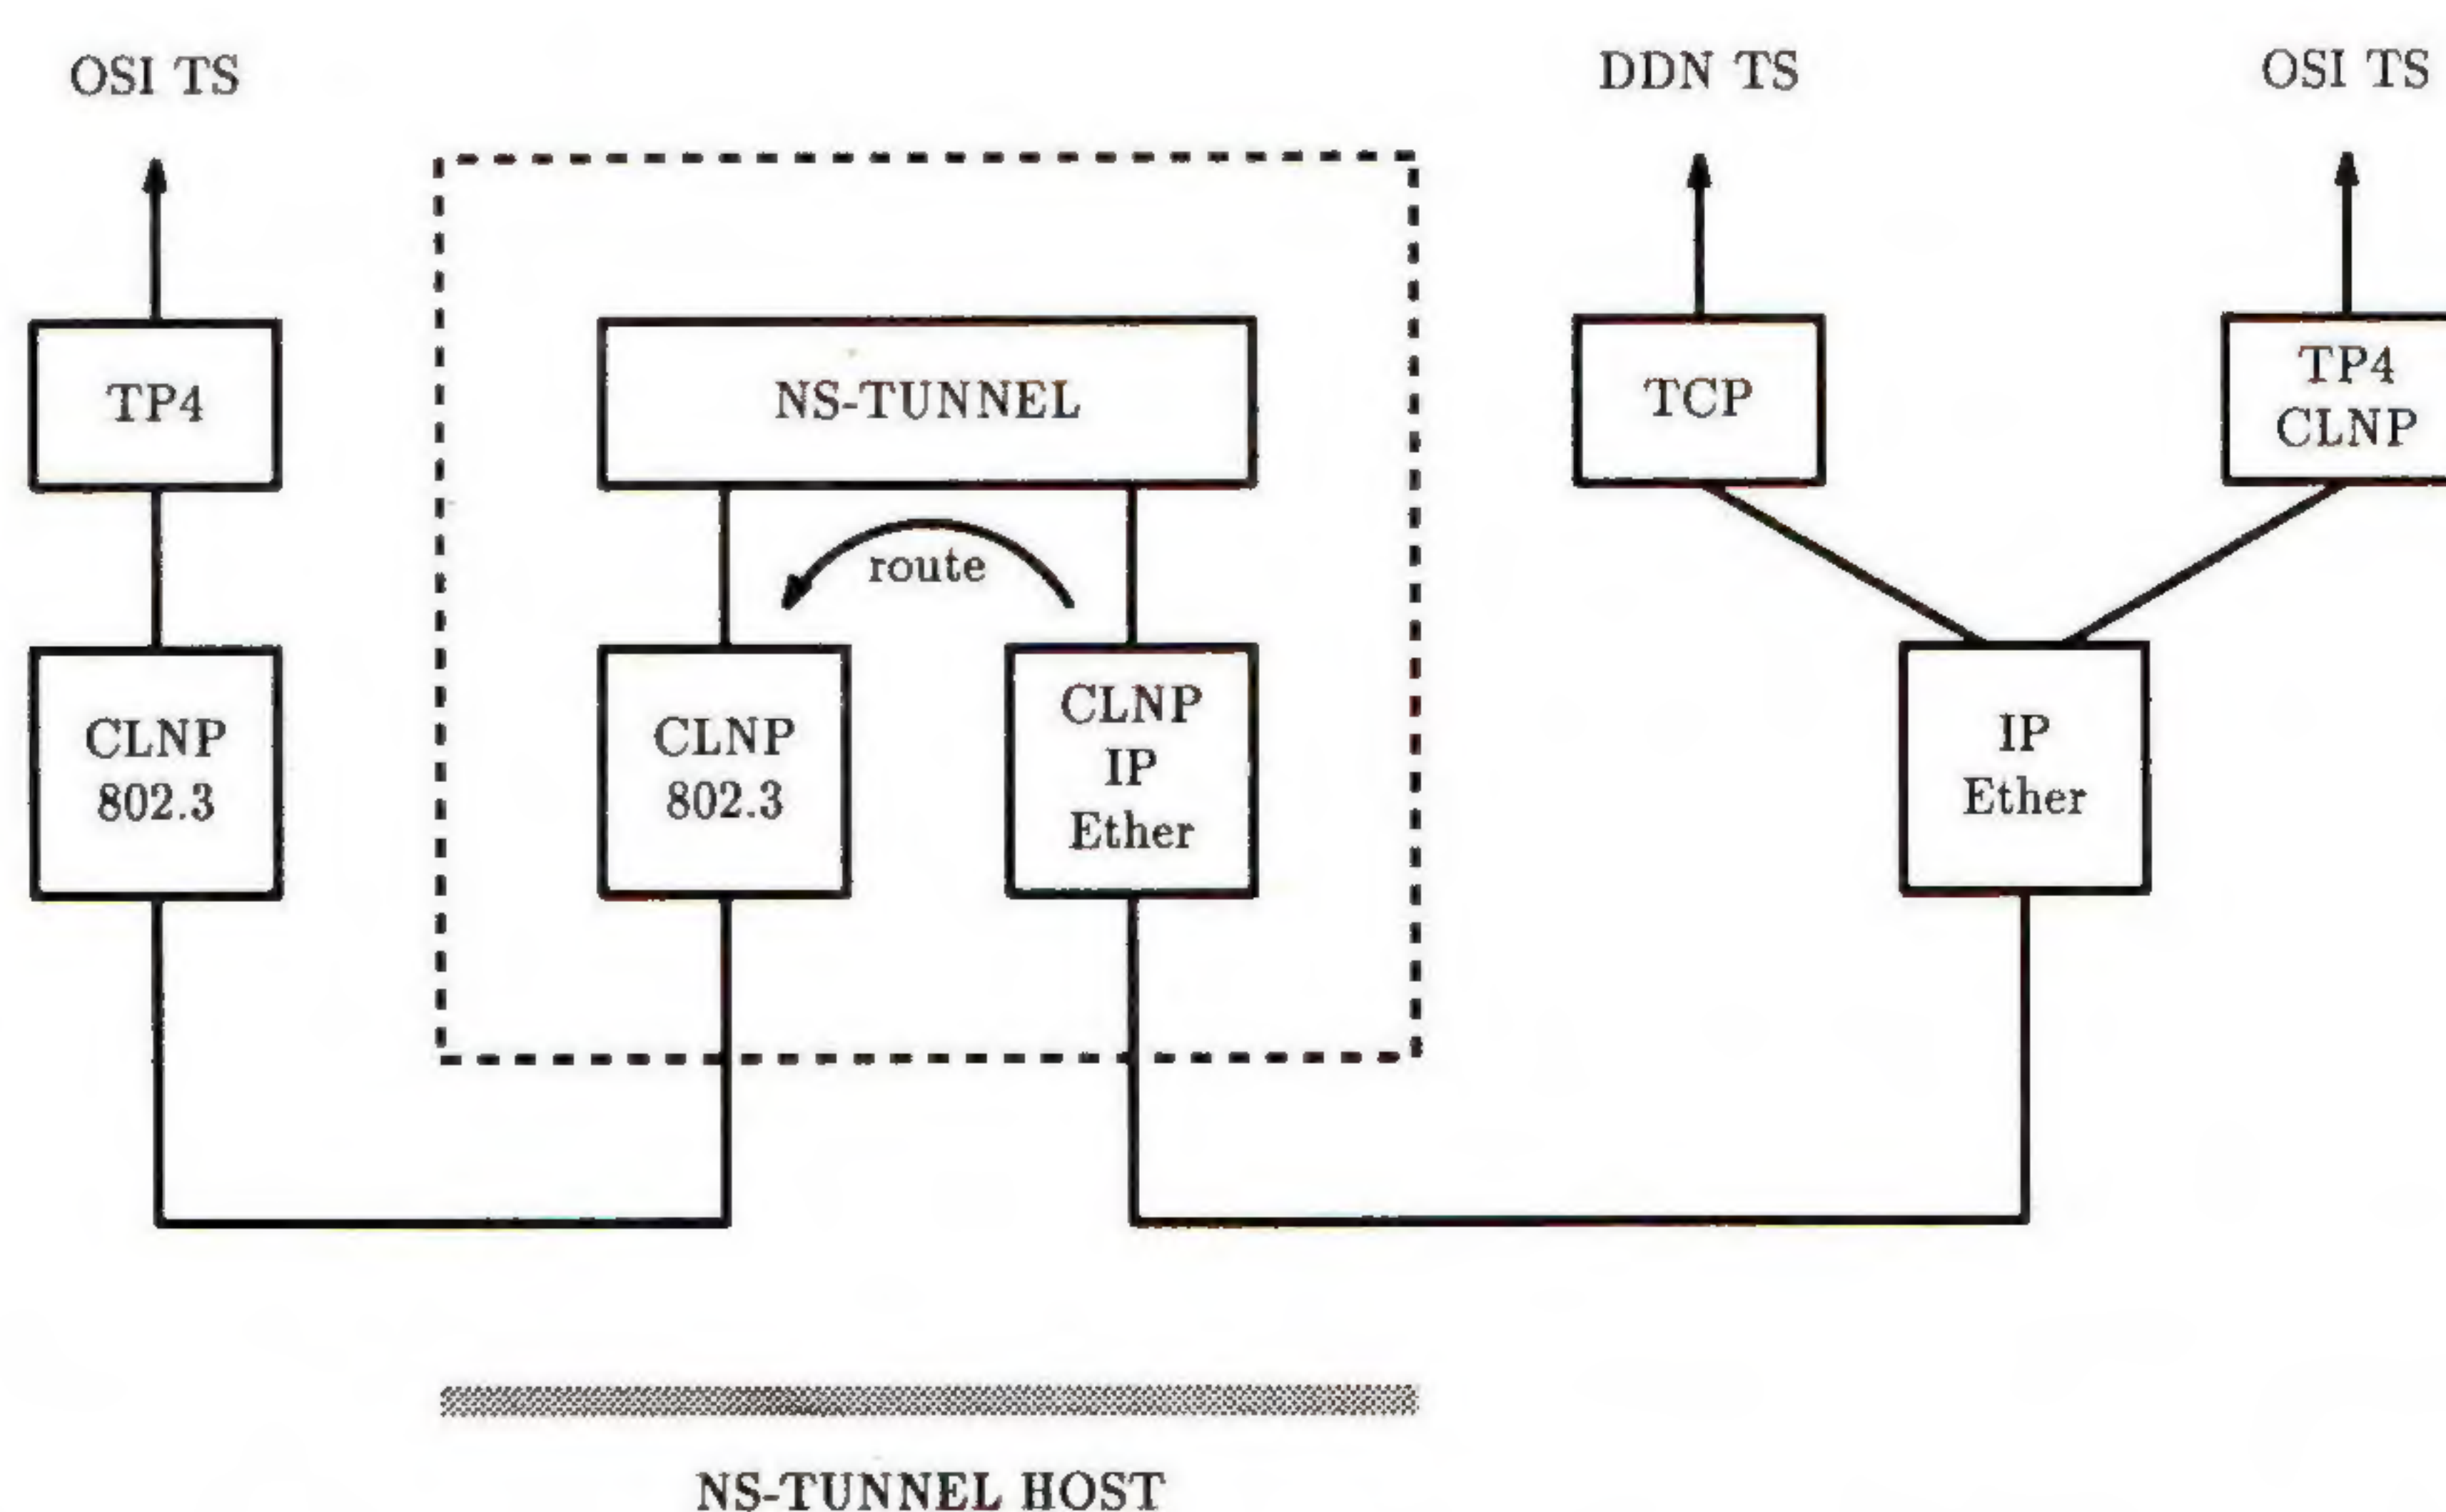


Figure 2: Network-Service Tunnel

Scorecard In terms of the metrics defined earlier:

- *Performance*: No worse than a typical CLNP-based router.
- *Flexibility*: High, as the NS-Tunnel is independent of any application.
- *Transparency*: Total, given adequate routing protocols.
- *Pervasiveness*: Participating TCP-based end-systems must now run both transport protocols (usually requiring "kernel" modifications on systems).

Other approaches There are several other approaches to the transition problem; this article has addressed only two. Some of the other approaches to be discussed in the tutorial included: dual-stacks, transport gateways, transport-service bridges [Ed.: See *ConneXions* Volume 2, No. 1, January, 1988], and the use of programmatic interfaces, such as AT&T's Transport Library Interface (TLI).

Conclusion None of the approaches presented will be found to have uniformly high scores across the metrics chosen for comparison. However, against a subset of the metrics, each approach can be seen to have advantages over the others. As the tutorial draws to a close, a strategy will be suggested that leverages the strengths of a few of the approaches in order to provide for a "minimal pain" transition plan. That strategy is.... That would be telling. Come to the tutorial and find out for yourself!

[Ed.: For those who missed the OSI Product Integration Conference, these issues will also be discussed in Advanced Computing Environments' April tutorials, scheduled for April 3—6 in Boston, MA. See also page 15].

References [Comer] *Internetworking With TCP/IP—Principles, Protocols, and Architecture*, Douglas E. Comer, Prentice-Hall, ISBN 0-13-470154-2.

[Tannenbaum] *Computer Networks*, Second Edition, Andrew S. Tannenbaum, Prentice-Hall, ISBN 0-13-162959-X.

MARSHALL T. ROSE is a Principal Software Engineer at The Wollongong Group, Inc., where he works on OSI protocols and transition strategies. He is the principal implementor of the ISO Development Environment (ISODE), an openly available implementation of the upper layers of the OSI protocol suite. He was co-author of RFC 1006 ("ISO Transport Services on top of the TCP"), and was a member of the IFIP working group committee whose efforts led to RFC 987 ("Mapping between X.400 and RFC 822"). He was co-author of RFC 1065 ("Structure and Identification of Management Information for TCP/IP-based internets") and RFC 1066 ("Management Information Base for Network Management of TCP/IP-based internets"). He also chaired the SNMP Extensions Working Group of the IETF which produced RFC 1067 ("A Simple Network Management Protocol"). These three documents form the short-term network management solution for TCP/IP-based networks. He is currently an advisor to the National Science Foundation, serving on its Network Technical Advisory Group. Rose received the Ph.D. degree in Information and Computer Science from the University of California, Irvine, in 1984.

The new NSFNET backbone network

by Hans-Werner Braun, Merit Computer Network
and University of Michigan

Introduction With the second generation of the National Science Foundation's backbone network infrastructure, July 1988 marked the beginning of national networking at megabit speeds. During July the older 56Kbps NSFNET backbone network with its six packet switching nodes (the so-called "Fuzzballs") was completely phased out and replaced by a new thirteen node backbone connected by T1 links at 1.544Mbps.

Funding The proposal to replace the old NSFNET backbone was submitted in August 1987 by MERIT, Inc., a computer network consortium of eight state-supported universities in Michigan, following a solicitation for proposals by the National Science Foundation. MERIT undertook this endeavor jointly with IBM Corporation and MCI Telecommunications Corporation. Further financial support is provided by the state of Michigan on the order of \$5 Million during the five year time frame of the award. In November 1987 Merit received the \$14M five year award, as a cooperative agreement with the NSF, in response to the proposal which outlined an operational new network for July 1988. The implementation of the network on schedule and within budget emphasized that cooperation between the federal government, industry, and universities can bear results of national interest to move into a leadership of collaborative high speed networking.

Three-level hierarchy The NSFNET model incorporates a three level hierarchy of campus networks, mid-level infrastructure and the NSFNET backbone itself into a giant structure of a meta-network connecting hundreds of individual subnetworks. The backbone allows for peer network connectivity to networks like national backbones (e.g., the NSN and the Arpanet) as well as international connections to national backbones of foreign countries.

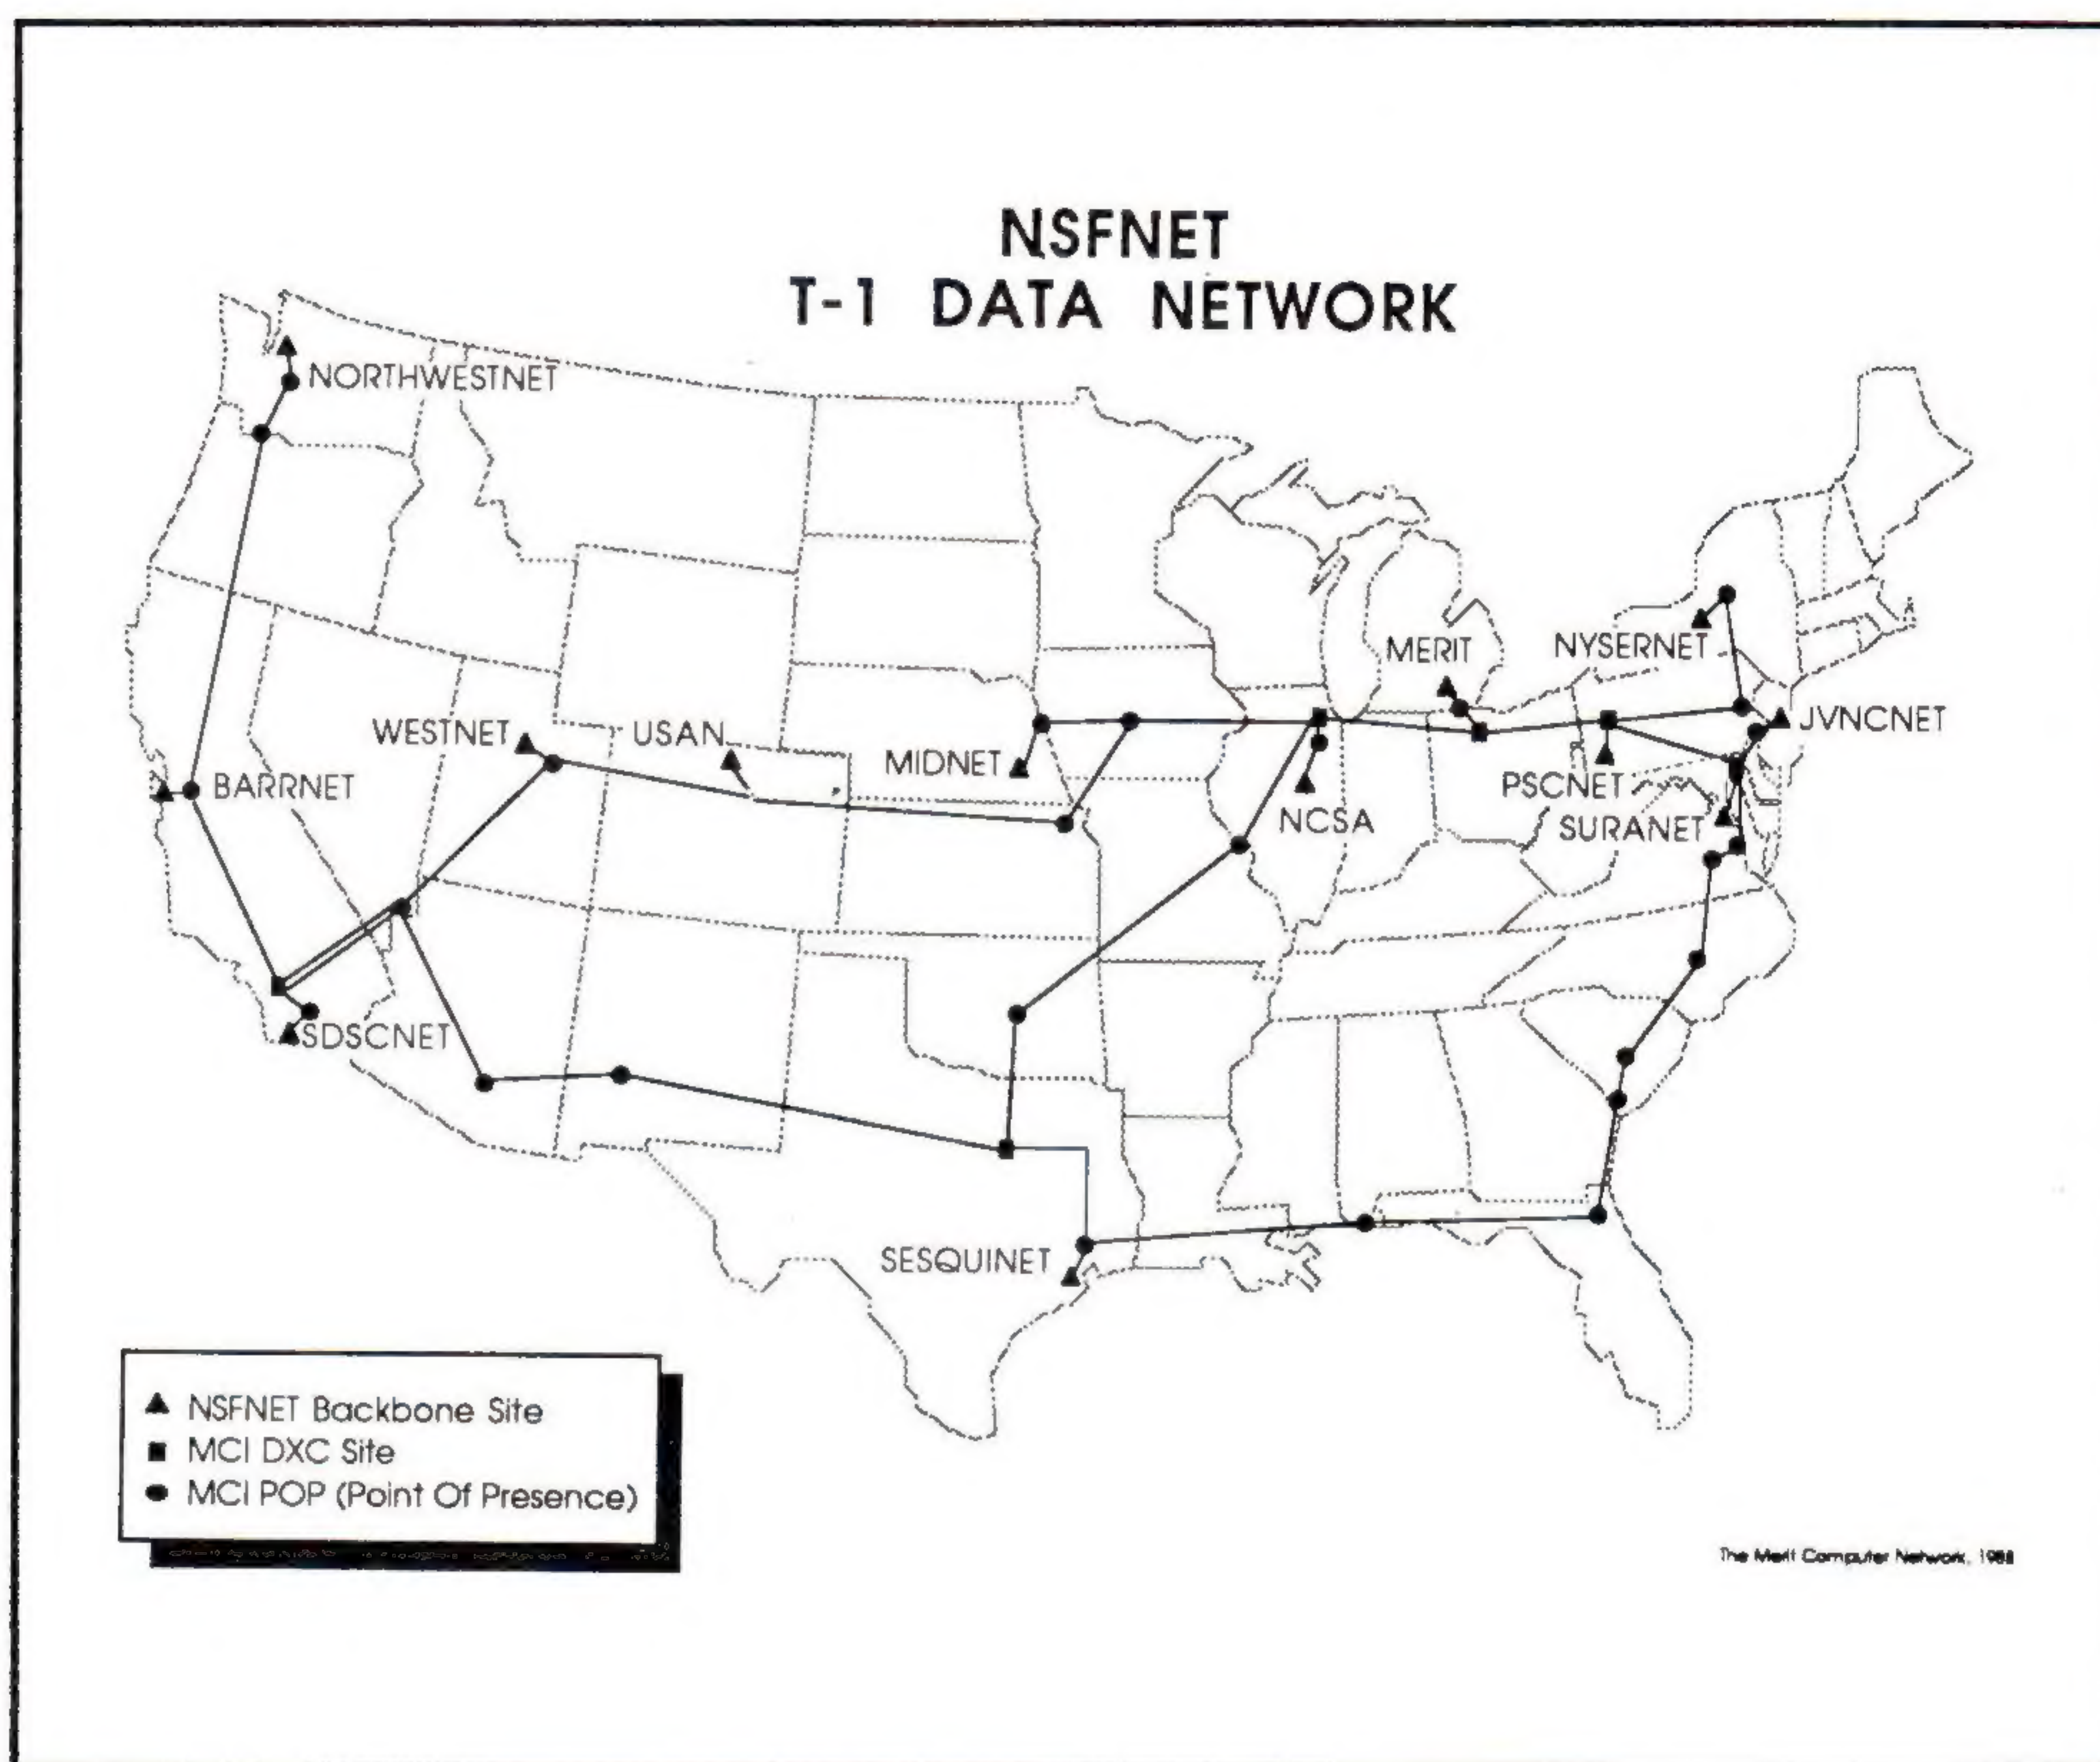
The NSFNET initially connects to thirteen mid-level networks covering geographic areas within the continental United States:

- BARRNET (Bay Area Regional Research Network)
- JVNCNET (John Von Neumann Center regional Network)
- Merit Computer Network
- MIDNET (Midwest Network)
- NCSA (National Center for Supercomputer Applications)
- NORTHWESTNET (Northwestern States Network)
- NYSERNet (New York State Education and Research Network) and the CNSF (Cornell National Supercomputing Facility)
- PSCNET (Pittsburgh Supercomputer Center Network)
- SDSCNET (San Diego Supercomputer Center Network)
- SESQUINET (Texas Sesquicentennial Network)

- USAN (University Satellite Network) and NCAR (National Center for Atmospheric Research)
- WESTNET (Southwestern States Network)

Backbone

The NSFNET Wide Area Communication Subsystem (WACS) is supplied by MCI and currently routes dedicated T1 circuits into the thirteen sites. The plan is to allow for digital reconfiguration within the MCI network as early as 1989, and for extensive access to MCI monitoring facilities within the same time frame. The access to sophisticated link level monitoring will be augmented with the use of the Extended Superframe Format (ESF) in the near future. The routing of the initial backbone circuits is illustrated below.



IDNX

In the initial implementation, the T1 circuits are terminated by Verilink Channel Service Units (CSU) which then connect to an Integrated Digital Network Exchange (IDNX). The IDNX instruments the ability to do dynamic allocation of bandwidth and routes within the physical limitations of the T1 circuits. This allows for the creation of a logical topology of sub-T1 channels on top of the physical structure provided by the underlying link level structure. The initial logical topology of this hybrid circuit/packet switching network is described in the diagram on the following page.


All the logical connections terminate in RS422 interfaces to processors in a Nodal Switching Subsystem (NSS), the packet-switching nodes of the NSFNET backbone. Each NSS is a loosely coupled cluster of RISC processor systems connected by a set of dual Token Rings. IBM RT/PCs are used for the RISC processor systems. Each packet-forwarder, be it a Packet Switching Processor (PSP) to one of the synchronous links or an Exterior Packet Switching Processor (E-PSP) connecting to a mid-level network, consists of a separate RISC processor system.

continued on next page

The new NSFNET backbone network (continued)

Routing

Another dedicated system, the Routing Control Processor (RCP), is responsible for routing computations and routing table management within the NSS. The RCP floods the PSPs and uses the ANSI IS-IS routing protocol for its inter-NSS routing while communicating with the remote NSS. The Exterior Gateway Protocol (EGP) is used for communication between an RCP and the appropriate gateway at the mid-level network. There are further individual processor systems in each NSS for network monitoring functions.



Network Management

Besides supplying the NSS, IDNX and CSU equipment, IBM is supplying extensive network management tools. Part of the network management is using the IBM NetView™, which is running on one of the two IBM 4381 mainframes contributed by IBM to this project. The NetView functionality is furthermore augmented by Internet standard network management protocols and other monitoring systems.

Project Management

Merit's NSFNET Project Director is Eric Aupperle, who is also the Director of the Merit Computer Network. Hans-Werner Braun is the Principal Investigator on this award. An Executive Committee (chaired by Douglas E. Van Houweling) and a Technical Committee (chaired by Hans-Werner Braun) aid the project. The partners on the project are jointly represented in these committees.

Organizationally, three groups have various project responsibilities:

- Internet Engineering, headed by Hans-Werner Braun
- Information Services, headed by James Sweeton
- Network Operations and Management, or the NOC, headed by Dale Johnson

The on-site IBM project management staff under IBM's Jack Drescher as well as the local MCI management team are dedicated and committed to the success of the NSFNET. They furthermore aid towards responsiveness in case of any issues and in accomplishing network growth as well as functionality improvements.

Future

During the five year time frame it is expected that this project will provide support for ISO protocols and, optionally, implement an upgrade to a full T3 network.

Information about the NSFNET project can be obtained from:

Merit Computer Network
1075 Beal Avenue
Ann Arbor, Michigan 48109-2112
(800) 66-MERIT
NSFNET-Info@MERIT.EDU

HANS-WERNER BRAUN received his engineering degree in 1978 in West Germany, following which he worked for five years in network engineering at the Regional Computing Center of the University of Cologne. Since April 1983 he has been at the University of Michigan's Computing Center and the Merit Computer Network, working on a variety of networking related projects. He chaired the Technical Committee of the National Science Foundation's Network Program Advisory Group (NPAG-TC) from February until December 1987 (when the NPAG got resolved) and is a member of the Network Technical Advisory Committee (NTAC) of the John von Neumann National Supercomputer Center (JvNC). He participates in the Internet Engineering Task Force and its steering group as well as in the Internet Architecture Task Force. He participated in several meetings for the design, implementation and operation of the first NSFNET backbone as well as for various mid-level networks attached to the backbone. Hans-Werner is Principal Investigator on the NSF project of the "Management and Operation of the NSFNET Backbone Network" and chairman of its joint Network Technical Committee. He was also Principal Investigator on a NSF grant to coordinate routing in the days of the old NSFNET backbone.

Background File Transfer Program (BFTP)

by Annette L. DeSchon and Robert T. Braden,
USC Information Sciences Institute

Introduction

The USC Information Sciences Institute has recently created a background file transfer service for use within the Internet.

Previously, file transfer in the Internet was usually implemented as an interactive, or "foreground" service. Users of a foreground type of service run a local FTP user interface program and request that a file transfer occur in real time. The user waits for the transfer to complete, and if the transfer fails at any time during the process, the user must issue another request.

More recently, as the Internet has grown, it has become increasingly subject to congestion and long delays, particularly during times of peak usage. In addition, planned and unplanned outages of hosts, gateways, and networks sometimes make it difficult for users to successfully transfer files in foreground. Performing file transfer asynchronously in the "background," provides a solution to some of these problems, by eliminating the requirement for a human user to be directly involved at the time that a file transfer takes place.

The BFTP server is built upon the third-party transfer model of FTP, described in RFC 959. Since BFTP coordinates file transfers between existing FTP server implementations, no changes to existing Internet protocol standards are required. This article presents a summary of the capabilities of BFTP, and is based on RFC 1068.

Potential advantages

Background file transfer has a number of potential advantages for the user:

- *No Waiting:* The user can request a large transfer and ignore it until a notification message arrives through some common channel, such as electronic mail.
- *End-to-end Reliability:* BFTP can try a transfer repeatedly until it either succeeds or fails permanently. This provides reliable end-to-end delivery of a file, in spite of the source or destination host being down, or poor Internet connectivity during some time period.
- *Deferred Delivery:* The user may wish to defer a large transfer until an off-peak period. This may become important when parts of the Internet adopt accounting and traffic-based cost-recovery mechanisms.

File transfer mechanisms

A background file transfer service requires two components: a user interface program to collect the parameters describing the required transfer(s), and a *file transfer control* (FTC) daemon to carry them out. In the BFTP design, the user interface program and the FTC daemon program execute on the same host, which we call the *BFTP control host*.

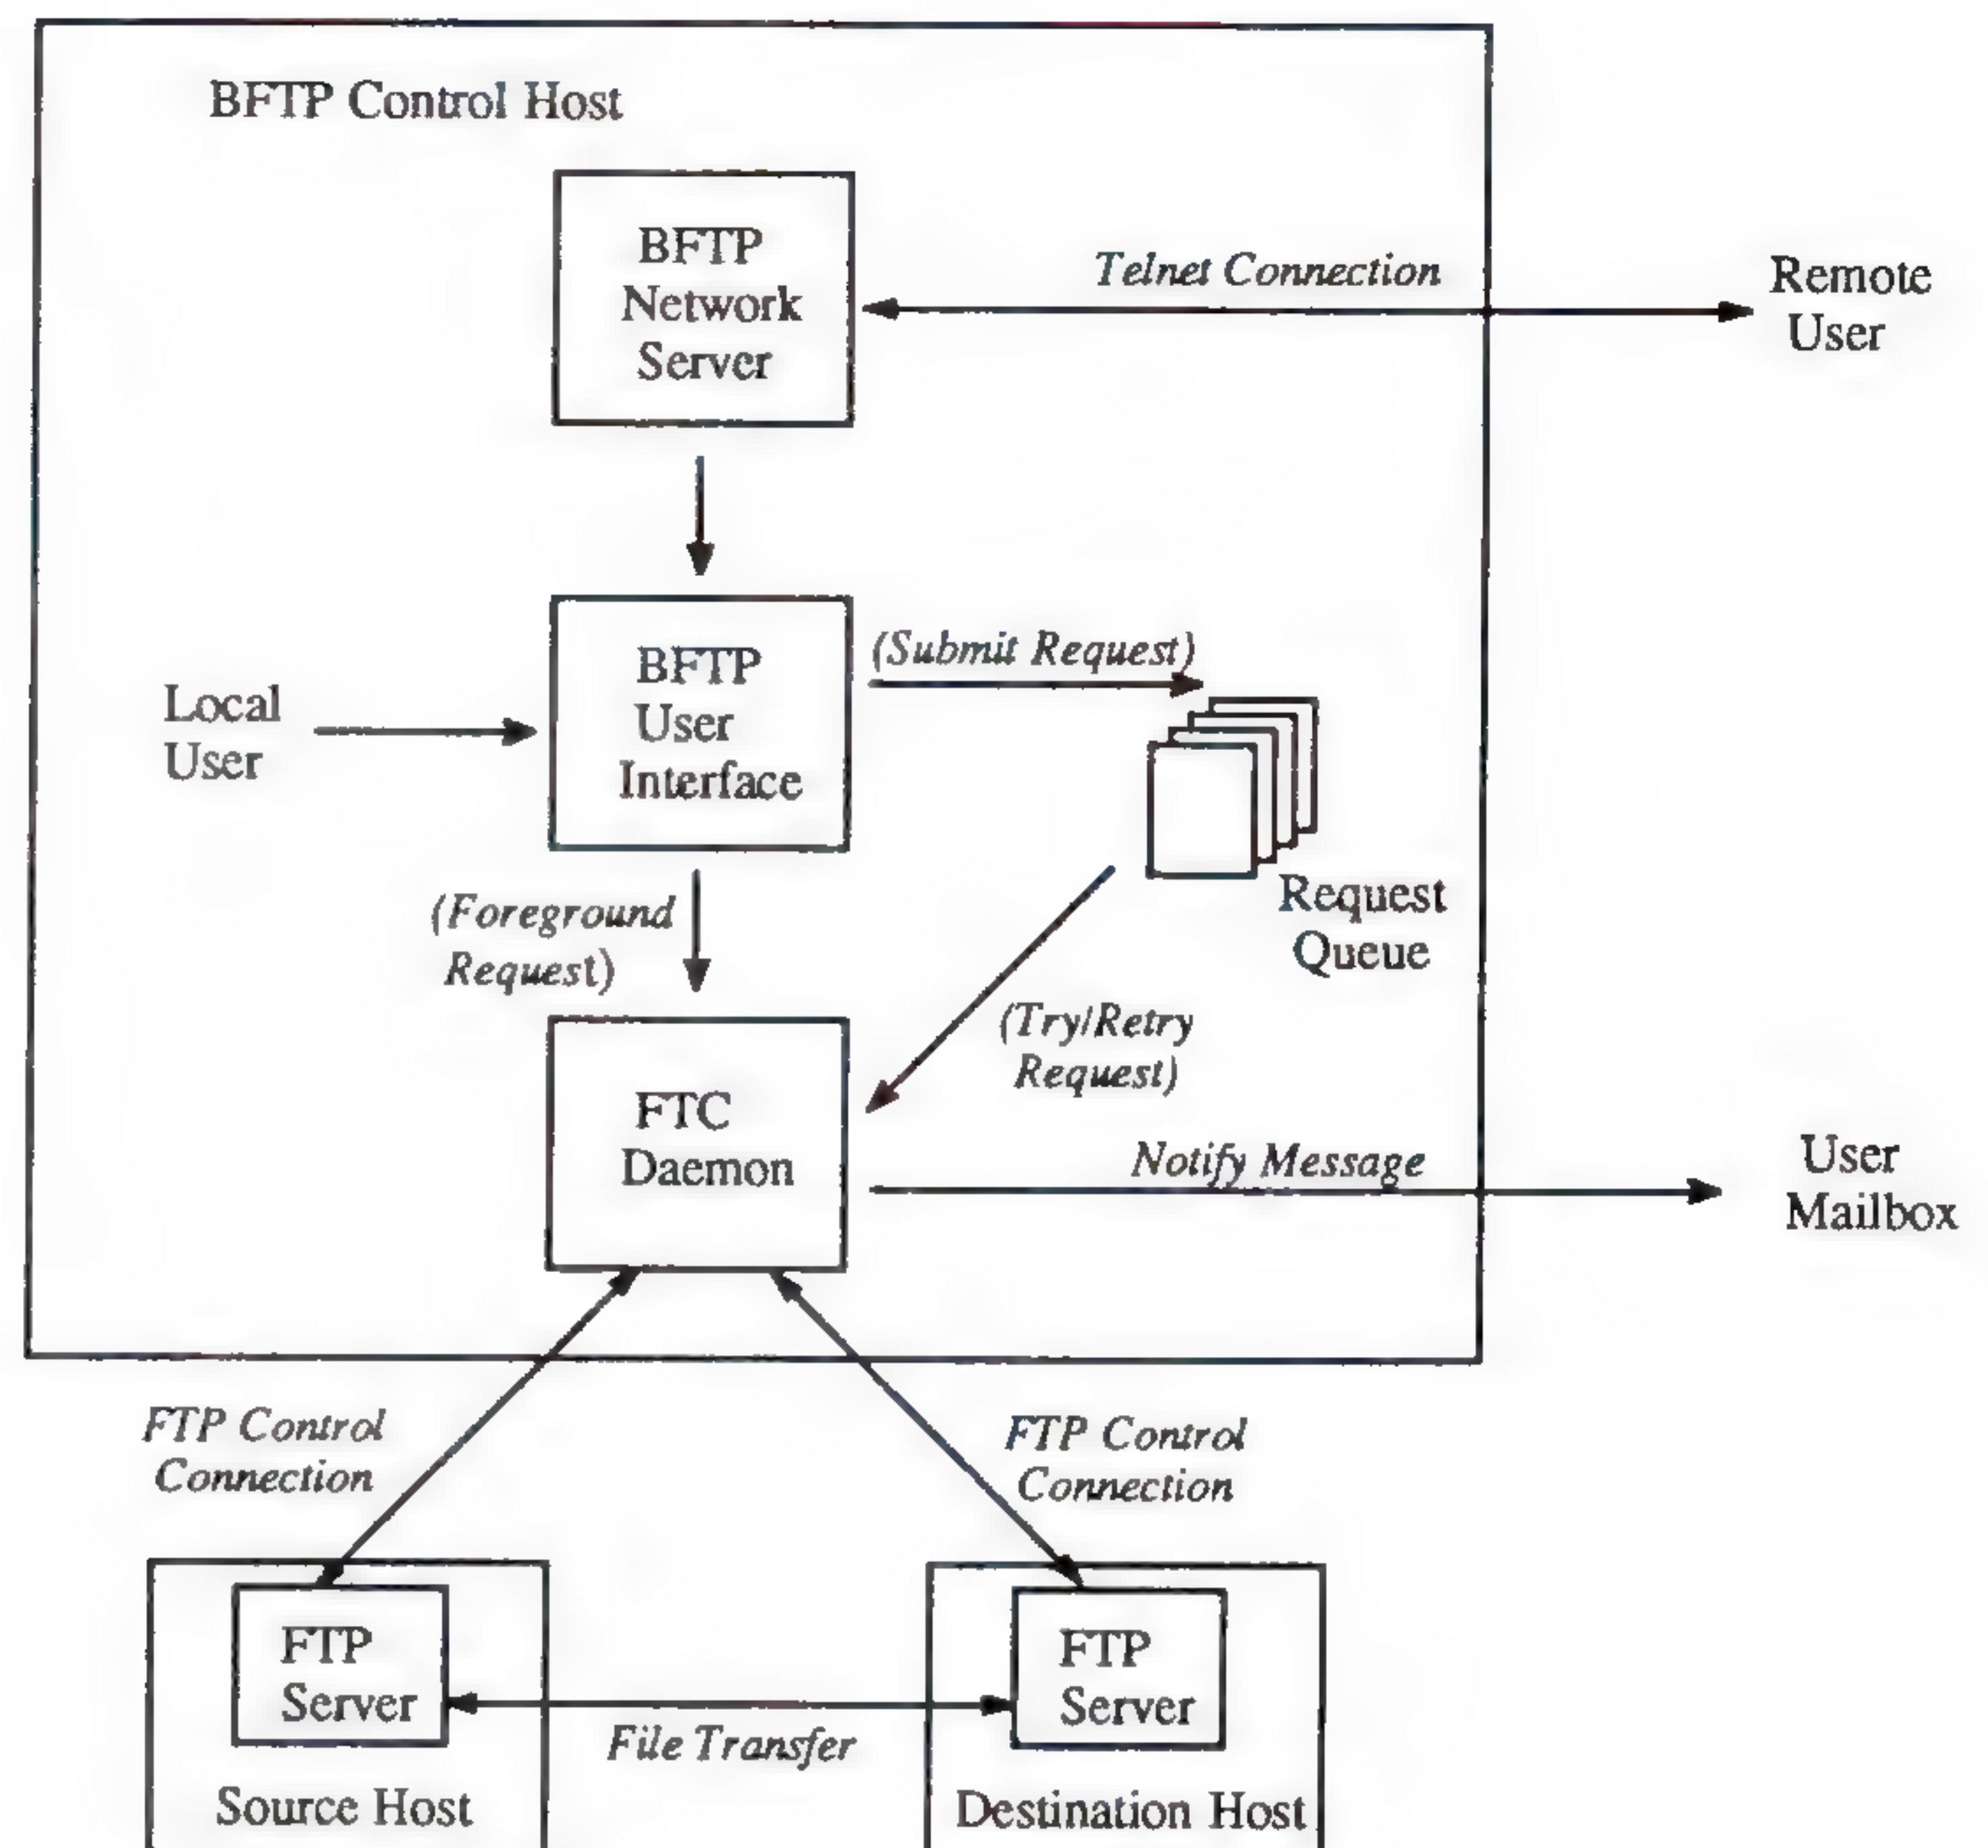


Figure 1: BFTP Model of operation

User interface features

The purpose of BFTP is to simplify the file transfer process and to place the burden of reliability on the BFTP control host. We have attempted to provide a "user friendly" command interface to BFTP, similar in flavor to the user interface of the TOPS-20 operating system. This interface provides extensive prompting, defaulting, and help facilities for every command.

We have also implemented a window-based BFTP user interface, the *bftptool*, which runs on a Sun workstation, in the *SunView* environment. (See Figure 2 on the following page) The various file transfer parameters appear in a form-style interface; defaults and multiple-choice style parameter values can be filled in via menus. An advantage of this form-style interface program is that it is possible to view all of the file transfer parameters simultaneously, providing the user with a sense of which parameter values might be mutually exclusive. Help information for each parameter and for each command is available via a mouse operation.

Reliable delivery

The reliable delivery function of BFTP, which is available via the "submit" command, is analogous to reliable delivery in a transport protocol like TCP. Both depend upon repeated delivery attempts until success is achieved, and in both cases the choice of the retry interval requires some care to balance overhead against unresponsiveness.

Human beings are impatient, but their impatience has a limit. If a file cannot be transferred "soon," a human being will turn to another project; typically, there is a tendency for the transfer to become less urgent the longer the wait. The FTC daemon of BFTP therefore starts each transfer request with a short default retry interval, e.g. 15 minutes, and then doubles this interval for successive retries, until a maximum interval, e.g. 4 hours, is reached. This is essentially the exponential backoff algorithm of the Ethernet, which is also used by transport protocols such as TCP, although BFTP and TCP have quite different rationales for the algorithm.

Background File Transfer Program (continued)

BFTP Tool

Quit

Explain

Verify

Submit

Transfer

Find

Clear

LoadReq

StoreReq

PurgeReq

RequestName: nicTest

Operation: copy Verbose: False

Source

Host: sri-nic.arpa

Dir: netinfo

File: nic-pubs.txt

Login: anonymous

Acct:

Password:

Port: 21

Destination

Host: venera.isi.edu

Dir:

File: nic-pubs.txt

Login: deschon

Acct:

Password:

Port: 21

FTP Parameters

Structure: file

Mode: stream

Type: ascii

Format: nonprint

MultipleMatching: FALSE

AppendFile: FALSE

UniqueFileName: FALSE

Introduction:

This Background File Transfer program may be used to submit a request to have a file transferred at some time in the future.

Mouse Functions:

The LEFT mouse button is used to position the cursor in a text field, and to select a command.

The CENTER mouse button is used to display help information in the text subwindow of the tool. To test this feature, position the mouse pointer over a command/field label, such as the RequestName, and click.

The RIGHT mouse button is used to display and fill in default values. To test this feature, position the mouse pointer over the label for a field, such as the Host field. Depress the right mouse button, highlight a value on the menu, and release the mouse button.

Figure 2: BFTP Tool

Verification

A potential human-engineering problem is that if the user makes a mistake in entering parameters, his mistake may not become apparent until much later. To avoid this problem, the BFTP user interface programs provide a command to “verify” the correctness of as many of the parameters as possible. Of course, such foreground verification of parameters is not possible if the remote host to which the parameters apply is currently unreachable.

Experience

BFTP has been available to users at ISI for some months. Users have reported a number of advantages of using BFTP:

- Some users prefer the prompting style of BFTP to the user interface of the foreground FTP they normally use.
- The BFTP “verify” command allows the user to verify that host names, passwords, and filenames are correct without having to wait for the entire transfer to take place.
- Since results are returned through the mail system, a transfer can occur without tying up a terminal line, a phone line, or even a window.

BFTP must be able to communicate with a variety of Server-FTP implementations, and we have observed much variation in the commands supported, error handling, and the timing in these servers. To diagnose problems that do occur, we have found it very useful to have a complete record of the interchange between the FTC daemon and the two FTP servers. This record is saved and is always included in the notification message mailed to the user.

Lastly, BFTP would benefit from the resolution of the following open protocol issues:

- There currently exist no provisions for Internet-wide user authentication. In the BFTP context, this means that passwords required for a file transfer must be present in BFTP request files. The security of these passwords is subject to the limitations of the file system security on the BFTP control host. Anonymous file transfer provides a partial solution, but a more general, long term solution is needed.
- Better mechanisms are needed to cope with the diversity of real file systems in the Internet.

For example, an extension could be made to the FTP protocol to allow the daemon to learn the delimiter conventions of each host file system. This could allow a more flexible and powerful multiple-file facility in BFTP. This could include the automatic transfer of directory subtrees, for example.

For more information on BFTP, contact:

Annette DeSchon
USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, Ca. 90292
213-822-1511 or deschon@isi.edu

References

DeSchon, A., & Braden, R., "Background File Transfer Program (BFTP)," RFC 1068, August 1988.

Postel, J., "File Transfer Protocol (FTP)," RFC 959, October 1985.

BFTP development was supported by the National Science Foundation under contract NCR-8718217.

ANNETTE DeSCHON is a member of the research staff at the USC Information Sciences Institute. Her interests include network protocols, internet-working and electronic mail. She received her B.A. from the University of California, Los Angeles, in 1975.

BOB BRADEN has spent nearly a lifetime in the computer field, starting with 0th generation computers in 1951. In 1986 he came to ISI from UCLA, where he spent the preceding 18 years working for the computer center. He first became involved with the Arpanet in 1970, when his systems group made the UCLA 360/91 an Arpanet host (1/1, now 10.1.0.1!). He participated in design of the Arpanet FTP and RJE protocols and in TCP development, and wrote the UCLA MVS TCP/IP code. He is a charter member of the IAB and is chairman of the End-to-End Task Force. Bob definitely qualifies as an *Internet Old Boy*.

The Users' Directory of Computer Networks

by Tracy L. LaQuey, The University of Texas System
Office of Telecommunication Services

Introduction

The "Users' Directory of Computer Networks" is a detailed and comprehensive compilation of host level information about various academic and research computer networks. Available in hard copy or on-line formats, this document provides general and/or host information for such networks as Arpanet/MILNET, BITNET, CSNET, ESnet, NSFNET, Public X.25 Networks, SPAN, THEnet and USENET. It also includes an electronic mail tutorial, a chapter on Domain Names, an index of hosts from all the networks grouped by organization, and network maps (hard copy version only). The directory reveals a "metanetwork" (a term coined by John Quarterman) that is accessible by hosts in the Texas Higher Education Network, (THEnet) explaining how to communicate through links between different national and international networks. It also puts into perspective the magnitude of this "metanetwork." This work has been in print for almost two years and has proven useful to all types of users, including those outside of THEnet, by serving as a reference guide and an educational tool.

Updates

A problem with compiling this type of directory is the volatility of the host and network information: even network information centers have difficulty keeping up with this task. Obviously, a hard-copy directory will not be totally accurate. However, like the telephone directory, which is somewhat out of date by the time it is printed, the "Users' Directory" furnishes readily accessible information, reduces support tasks and encourages people to keep entries up-to-date. It provides a good first step toward finding information if other methods are not available. Currently, the hard-copy document will be updated and reprinted each July, and the on-line version will be revised each January and July.

History

William C. Bard, Director of The University of Texas System Office of Telecommunication Services, conceived the idea for a definitive directory for academic and research networks in early 1987. The current edition, compiled by myself, was based on the first edition compiled and written by Carol Engelhardt Kroll. The objective was to reduce the number of queries from people wanting information about various networks, desiring to contact people, and wanting to find out where certain resources were available. Most of the information was obtained from various network information centers. One of the more helpful documents for developing the first edition was "Notable Computer Networks," CACM, Oct. 1986, by Quarterman and Hoskins.

Ordering information

To order the directory, send \$17 to:

The University of Texas System
Office of Telecommunication Services
Balcones Research Center
10100 Burnet Road
Austin, TX 78758-4497 netbook@thenic.the.net

The on-line version can be obtained by anonymous Internet FTP from: emx.utexas.edu (128.83.1.33) in the directory "net.directory." The file "ftp.instructions" describes each section and also contains information for ordering hard copies.

Looking ahead to 1989

Included with this month's *ConneXions* is a Table of Contents for Volume 2, 1988 which provides a quick-reference to past articles. (A similar table is available for Volume 1, 1987). All back issues can be ordered for \$10 each. Next year we plan to produce a special *ConneXions* binder, more on that in a future issue.

In 1989, we will cover a number of topics related to protocol standardization, interoperability, and transition to OSI. Some highlights include:

- A series of articles on the various *Components of OSI*.
- A special issue on *Network Management* with details on both the SNMP and CMOT architectures.
- A special issue on *Subnets*.
- An article on *Interoperability beyond TCP/IP*, what is happening with databases and distributed systems?
- A full report on the *virus* or *worm* which attacked the Internet in early November, 1988.

We hope you have enjoyed this volume. As always, we welcome your comments and suggestions.

Tutorials

Next on our calendar is a number of 2-day tutorials, which will be held at the Lafayette Hotel in Boston April 3—6, 1989. The program is as follows:

Monday and Tuesday:

- | | |
|--|----------------|
| • <i>Introduction to TCP/IP</i> | Doug Comer |
| • <i>Berkeley UNIX Networking</i> | Mike Karels |
| • <i>TCP/IP for the VM Systems Programmer</i> | Mike Hojnowski |
| • <i>Local Area Networks
and TCP/IP Alternatives</i> | Charles Brown |
| • <i>System Administration</i> | Eric Brunner |

Wednesday and Thursday:

- | | |
|--|--------------------------------|
| • <i>The Domain Name System</i> | Paul Mockapetris |
| • <i>Network Management
of TCP/IP-based Internets</i> | Jeff Case |
| • <i>Bridges and Routers</i> | Radia Perlman |
| • <i>Practical Perspectives
on OSI Networking</i> | Marshall Rose &
Chris Moore |
| • <i>Message Handling Systems
and Directory Services—X.400/X.500</i> | Jim White &
Ted Myer |

A complete tutorial program will be mailed to you in the near future. For more information, call Advanced Computing Environments at 415-941-3399.

CONNEXIONS

480 San Antonio Road
Suite 100
Mountain View, CA 94040

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

CONNEXIONS

PUBLISHER Daniel C. Lynch

EDITOR Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President, National Research Initiatives.

Dr. David D. Clark, The Internet Architect, Massachusetts Institute of Technology.

Dr. David L. Mills, NSFnet Technical Advisor; Professor, University of Delaware.

Dr. Jonathan B. Postel, Assistant Internet Architect, Internet Activities Board; Associate Director, University of Southern California Information Sciences Institute.

Subscribe to CONNEXIONS

U.S./Canada \$100. for 12 issues/year \$180. for 24 issues/two years \$240. for 36 issues/three years
International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNEXIONS).

☐ Charge my ☐ Visa ☐ Master Card Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to:

CONNEXIONS

480 San Antonio Road Suite 100
Mountain View, CA 94040
415-941-3399

Back issues available upon request \$10./each
Volume discounts available upon request

CONNEXIONS